

Claims

What is claimed is:

1. A method to use in a node within a network comprising a transport layer protocol
5 providing end to end data transfer, for multicasting datagrams on a virtual ring, each
node on the virtual ring being logically connected according to the network transport
layer protocol to two and only two neighbor nodes through virtual connections, an
upstream neighbor node and a downstream neighbor node, said method comprising the
steps of:

10

- sending a virtual ring datagram to the downstream neighbor node on the virtual ring;
said virtual ring datagram comprising:
 - a virtual ring identifier;
 - means for identifying the node originator of the virtual ring datagram;
 - 15 • data;

when a datagram is received;

- identifying the received datagram;

20 if the received datagram is a token:

- identifying the virtual ring;
- checking whether the token is valid or not;
- if the token is valid, forwarding the token to the downstream neighbor node on the
identified virtual ring.

25

if the received datagram is a virtual ring datagram:

- identifying the virtual ring;
- checking the node originator of the received virtual ring datagram;

30 if the received virtual ring datagram has not been locally originated:

- processing data comprised in said virtual ring datagram;

- forwarding said virtual ring datagram to the downstream neighbor node on the identified virtual ring;

if the received virtual ring datagram has been locally originated:

- 5 • removing the virtual ring datagram from the virtual ring.

2. The method according to any one of the preceding claims wherein

- 10 • the step of identifying the virtual ring when a token is received, comprises the further step of:

- checking that the token has been sent by the upstream neighbor node on the identified virtual ring;

- 15 • the step of identifying the virtual ring when a virtual datagram is received, comprises the further step of:

- checking that the virtual ring datagram has been sent by the upstream neighbor node on the identified virtual ring.

20

3. The method according to any one of the preceding claims wherein:

- 25 • a node on the virtual ring is defined as being a virtual ring manager node;
- the token comprises a sequence number incremented each time the token is received by the virtual ring manager node;
 - the step of checking whether the token is valid or not, comprises the further step of:
 - checking whether or not the token sequence number has been incremented since the last reception.

- 30 4. The method according to any one of the preceding claims wherein the step of checking whether or not the token is valid, comprises the further step of:

- if the token is not valid, executing a recovery procedure.

5. The method according to any one of the preceding claims wherein:

5

- the step of forwarding the token to the downstream neighbor node on the identified virtual ring, comprises the further step of:

10

- starting a timer and waiting for the return of the token;
- executing a recovery procedure when the timer expires;

- the step of, receiving a token comprises the further step of:

- stopping the timer.

15

6. The method according to any one of the preceding claims wherein a node is:

- a computer system routing datagrams in the network, preferably a router, or
- a computer system exchanging datagrams on the network, preferably a client or a server.

20

7. The method, to use in a virtual ring manager node, according to any one of the preceding claims, comprising the preliminary steps of:

- generating a token;
- 25 • setting the token sequence number to an initial value;
- forwarding said token to the downstream neighbor node on the virtual ring;

the step of forwarding the token to the downstream neighbor node on the identified virtual ring, comprising the further step of :

30

- incrementing the token sequence number;

the step of, executing a recovery procedure when the timer expires, comprising the further step of:

- 5 • generating a new token;
- forwarding said token to the downstream neighbor node on the virtual ring.

8. The method according to any one of the preceding claims wherein the token is a datagram comprising:

10

- a header comprising:
 - a source address of the sending node; and
 - a destination address of the next node on the virtual ring;

15

- a header comprising:
 - a source port; and
 - a destination port;
 - means for identifying the datagram as being a token;
 - means for identifying the virtual ring;
 - a sequence number incremented each time the token is received by the virtual ring
- 20 manager node.

9. The method according to any one of the preceding claims wherein each virtual ring datagram circulating on the virtual ring comprises:

25

- a header comprising:
 - a source address of the sending node on the virtual ring, and
 - a destination address of the next node on the virtual ring;

30

- a header comprising:
 - a source port; and
 - a destination port;
- a virtual ring header comprising:

- means for identifying the virtual ring datagram on the virtual ring;
- means for identifying the virtual ring;
- means for identifying the node originator of the virtual datagram;
- data.

5

10. The method according to any one of the preceding claims wherein comprising further the step of:

- maintaining and updating the following information:
 - 10 • means for identifying the virtual ring;
 - the address of the upstream neighbor node;
 - the address of the downstream neighbor node;
 - the address of the virtual ring manager;
- optionally, the address of a backup virtual ring manager.

15

11. The method according to any one of the preceding claims comprising the preliminary step of joining the virtual ring, said step comprising the steps of:

- sending to a virtual ring node manager node previously defined on the virtual ring,
 - 20 an insertion request message comprising:
 - the address of the node;
 - means for identifying the virtual ring;
- receiving an insertion confirmation message from the virtual ring manager node comprising:
 - 25 • the address of an upstream neighbor node;
 - the address of a downstream neighbor node.

12. The method according to the preceding claim,

30 wherein the step of sending an insertion request message, comprises the further step of:

- starting an insertion timer;

wherein the step of receiving an insertion confirmation message, comprises the further
5 step of:

- stopping the insertion timer;

and wherein, if the insertion timer expires, said method comprises the further steps of:

10

- sending an insertion request message comprising
 - the address of the node;
 - means for identifying the virtual ring;to a backup ring manager node previously defined on the virtual ring;

15

- restarting the insertion timer;
- receiving an insertion confirmation message from the backup virtual ring manager comprising:
 - the address of an upstream neighbor node;
 - the address of a downstream neighbor node;

20

- stopping the insertion timer.

13. The method according to any one of the preceding claims comprising the further
step of leaving the virtual ring, said step comprising the steps of:

25

- sending to a virtual ring manager node previously defined on the virtual ring, a removal request message comprising:
 - the address of the upstream neighbor node;
 - the address of the downstream neighbor node;
 - the address of the node;

30

- receiving a removal confirmation message from the virtual ring manager.

14. The method according to the preceding claim,

wherein the step of sending a removal request message, comprises the further step of:

- 5 • starting a removal timer;

wherein the step of receiving a removal confirmation message, comprises the further step of:

- 10 • stopping the removal timer;

and wherein, if the insertion timer expires, said method comprises the further steps of:

- 15 • sending to a backup ring manager node previously defined on the virtual ring, a removal request message, comprising:
- the address of the upstream neighbor node;
 - the address of the downstream neighbor node;
 - the address of the node;
- restarting the removal timer;
- 20 • receiving a removal confirmation message from the backup virtual ring manager;
- stopping the removal timer.

15. The method according to any one of the preceding claims comprising the further steps of :

25

- receiving from a virtual ring manager node defined on the virtual ring, a change neighbor message comprising:
 - the address of a new upstream neighbor node; or/and
 - the address of a new downstream neighbor node;
- 30 • maintaining:
 - the address of the new upstream neighbor node; or/and

- the address of the new downstream neighbor node;
- sending to the virtual ring manager node a neighbor changed confirmation message.

16. The method, to use in a virtual ring manager node, according to any one of the
5 preceding claims comprising the further steps of :

- receiving an insertion request message comprising the IP address of a new node to insert into the virtual ring;
- sending to the downstream neighbor node, a change neighbor message comprising:
10 • the address of new node considered as the new upstream neighbor node of said downstream neighbor node;
- receiving a neighbor changed confirmation message;
- updating the address of the downstream neighbor node with the address of the new node;
- 15 • sending to the new node an insertion confirmation message comprising:
 - the address of the upstream neighbor node of the new node;
 - the address of the downstream neighbor node of the new node.

17. The method, to use in a virtual ring manager node, according to any one of the
20 preceding claims comprising the further steps of :

- receiving from a node on the virtual ring, a removal request message comprising:
 - the address of the upstream node of the node to remove;
 - the address of the downstream node of the node to remove;
 - 25 • the address of the node to remove;
- sending to the downstream neighbor node of the node to remove, a change neighbor message comprising:
 - the address of the upstream neighbor node of the node to remove;
- sending to the upstream neighbor node of the node to remove, a change neighbor
30 message comprising:
 - the address of the downstream neighbor node of the node to remove;

- receiving a neighbor changed confirmation message from the upstream neighbor node and from the downstream neighbor node of the node to remove;
- sending to the node to remove a removal confirmation message.

5 18. The method, to use in a virtual ring manager node, according to any one of the preceding claims comprising the further steps of:

- receiving a neighbor loss indication message indicating a loss of connection with an upstream neighbor node on the virtual ring; said neighbor loss indication message
10 comprising:
 - the address of the upstream neighbor node in failure;
 - the address of the node that has originated the neighbor loss indication message and detected the loss of connection with its upstream neighbor node;
- sending to the downstream neighbor node of the node in failure a change neighbor
15 message comprising:
 - the address of a new upstream neighbor node;
- receiving a neighbor changed message from the downstream neighbor node of the node in failure;
- receiving a neighbor loss indication message indicating a loss of connection with a
20 downstream neighbor node on the virtual ring; said neighbor loss indication message comprising:
 - the address of the downstream neighbor node in failure;
 - the address of the node that has originated the neighbor loss indication message and detected the loss of connection with its downstream neighbor node;
- sending to the upstream neighbor node of the node in failure a change neighbor
25 message comprising:
 - the address of a new downstream neighbor node;
- receiving a neighbor changed message from the upstream neighbor node of the node in failure.

30

19. The method according to any one of the preceding claims wherein:

- said network is a TCP/IP (Transmission Control Protocol/Internet Protocol) network,
- logical connections between nodes part of the virtual ring are TCP sessions;
- nodes are multicasting datagrams with their neighbor nodes, on the virtual ring,
- 5 through said TCP sessions;
- addresses are IP addresses.

20. The method according to any one of the preceding claims wherein messages exchanged between nodes for inserting a node or removing a node, are UDP (User
10 Datagram Protocol) messages.

21. A node comprising means adapted for carrying out the method according to any one of the preceding claims.

15 22. A computer program comprising instructions for carrying out the method according to any one of claims 1 to 20 when said computer program is executed on a node.